arXiv:1612.09121v1 [stat.ME] 29 Dec 2016

On Perfect Clustering of High Dimensional Data

Soham Sarkar[1] and Anil K. Ghosh[2]

Theoretical Statistics and Mathematics Unit
Indian Statistical Institute
203, B. T. Road, Kolkata 700108

December 30, 2016

Abstract

Popular clustering algorithms based on usual distance functions (e.g., Euclidean distance) often suffer in high dimension, low sample size (HDLSS) situations, where concentration of pairwise distances has adverse effects on their performance. In this article, we use a dissimilarity measure based on the data cloud, called MADD, which takes care of this problem. MADD uses the distance concentration property to its advantage and, as a result, clustering algorithms based on MADD usually perform better for high dimensional data. We also address the problem of estimating the number of clusters. This is a very challenging problem in cluster analysis, and several algorithms have been proposed for it. We show that many of these existing algorithms have superior performance in high dimensions when MADD is used instead of Euclidean distance. We also construct a new estimator based on penalized Dunn index and prove its consistency in HDLSS asymptotic regime, where the sample size remains fixed and the dimension grows to infinity. Several simulated and real data sets are analyzed to demonstrate the importance of MADD for cluster analysis of high dimensional data.

Keywords: Dunn index, hierarchical clustering, high dimensional consistency, k-means clustering, pairwise distances, Rand index.

[1]Corresponding author. Email: sohamsarkar1991@gmail.com
[2]Email: akghosh@isical.ac.in

1 Introduction

Let $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n \in \mathbb{R}^d$ be a sample of n unlabeled observations coming from different populations. The aim of cluster analysis is to divide this sample into several groups of 'similar' observations. In practice, one uses an appropriate measure of similarity (or, dissimilarity) between a pair of observations, and a clustering algorithm is developed based on that. When the measurement variables are all continuous, a popular choice for the dissimilarity index is the Euclidean distance or the squared Euclidean distance. Popular clustering algorithms like k-means, k-medoids and hierarchical clustering based on different linkage methods (see Hastie et al. 2009; Duda et al. 2012) generally use dissimilarity indices based on the Euclidean distance. Spectral clustering algorithms (see von Luxburg 2007) often use similarity index based on radial basis function, which is a decreasing function of the Euclidean distance. These clustering algorithms work well when the sample size is sufficiently large compared to the data dimension. But, like other nonparametric methods, they often perform poorly in high dimension, low sample size (HDLSS) situations.

To demonstrate this, let us consider a simple example (call it Example A), where we have observations from two d-dimensional normal distributions. One of these distributions has the mean vector $\mathbf{0}_d = (0, \ldots 0)'$ and the dispersion matrix $\sigma_1^2 \mathbf{\Sigma}_d$, while the other one has the mean vector $\boldsymbol{\mu}_d = (\mu_1, \ldots, \mu_d)'$ and the dispersion matrix $\sigma_2^2 \mathbf{\Sigma}_d$. Here $\mu_i = (-1)^{i+1}$ for $i = 1, \ldots, d$, and $\mathbf{\Sigma}_d = ((\sigma_{ij}))_{d \times d}$ is a block diagonal matrix with $\sigma_{ii} = 1$ for $i = 1, 2, \ldots, d$, $\sigma_{(2i-1)2i} = \sigma_{2i(2i-1)} = 0.98$ for $i = 1, 2, \ldots, \lfloor d/2 \rfloor$ ($\lfloor t \rfloor$ is the highest integer $\leq t$) and $\sigma_{ij} = 0$ otherwise. Taking $\sigma_1^2 = 0.5$ and $\sigma_2^2 = 2$, we generated 50 observations from each distribution. Figure 1(a) shows the central regions of these two distributions with coverage probabilities 0.25, 0.5, 0.75 and 0.9 when $d = 2$. We used the average linkage method (AvgL) as well as the k-means algorithm (kM) based on Euclidean distance to estimate two clusters in the sample consisting of 100 observations. For $i = 1, \ldots, n$, let $\mathcal{C}(\mathbf{x}_i)$ be the actual cluster label of $\mathbf{x}_i$ and $\delta(\mathbf{x}_i)$ be the cluster label assigned to $\mathbf{x}_i$ by a clustering algorithm δ. We measure

the performance of δ using the Rand index (see Rand 1971)

$$\mathcal{R}(\delta) = \frac{2}{n(n-1)} \sum_{1 \le i < j \le n} \mathbb{I}\Big[\mathbb{I}\{\delta(\mathbf{x}_i) = \delta(\mathbf{x}_j)\} + \mathbb{I}\{\mathcal{C}(\mathbf{x}_i) = \mathcal{C}(\mathbf{x}_j)\} = 1\Big], \tag{1}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function. Note that δ leads to perfect clustering if $\delta(\mathbf{x}_i) = \pi\{\mathcal{C}(\mathbf{x}_i)\}$ for all $i = 1, \ldots, n$ and a suitable permutation π of $\{1, \ldots, \max\{\mathcal{C}(\mathbf{x}_1), \ldots, \mathcal{C}(\mathbf{x}_n)\}\}$. In that case, we have $R(\delta) = 0$. Higher values of Rand index indicate more deviation from perfect clustering. We repeated our experiment 100 times, and the average Rand index of an algorithm over these 100 trials was computed for $d = 2^r$, with $r = 1, 2, \ldots, 11$. In this example, though the two distributions do not have disjoint supports, the separation between them is quite evident (see Figure 1(a)), and it increases with the dimension. So, for a good clustering algorithm, the Rand index is expected to shrink to 0 as the dimension increases. But, that was not the case for AvgL and kM. Both of them had miserable performance for all values of d (see Figure 2(a)). The spectral clustering algorithm proposed by Shi and Malik (2000) also had higher Rand index when a similarity measure based on radial basis function was used (see Figure 2(a)).

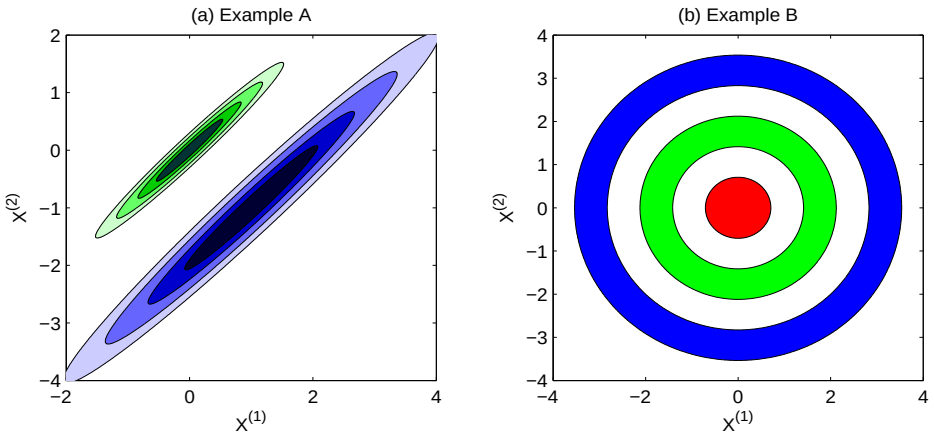

Figure 1: Different populations in Examples A and B when $d = 2$.

We carried out another experiment with observations generated from three distributions $U_d(0, 0.5)$, $U_d(1, 1.5)$ and $U_d(2, 2.5)$ with disjoint supports. Here $U_d(a, b)$ denotes a d-dimensional uniform distribution over the region $\{\mathbf{x} \in \mathbb{R}^d : a\sqrt{d} \le \|\mathbf{x}\| \le b\sqrt{d}\}$. Figure 1(b) shows the supports of these three distributions for $d = 2$. We generated 50 observations

from each distribution, and different clustering algorithms were used to divide these 150 observations into three different clusters. In this example (call it Example B) also, all these three methods, especially AvgL and kM, had poor performance (see Figure 2(b)).

Clustering algorithm based on maximal data pilling (MDP) distance (Ahn et al. 2012), which is especially designed for high dimensional data, performed well in Example A for $d \geq 2^7$, but it performed poorly for all smaller values of d. In Example B, its performance was even worse. It had much higher Rand index for all values of d.

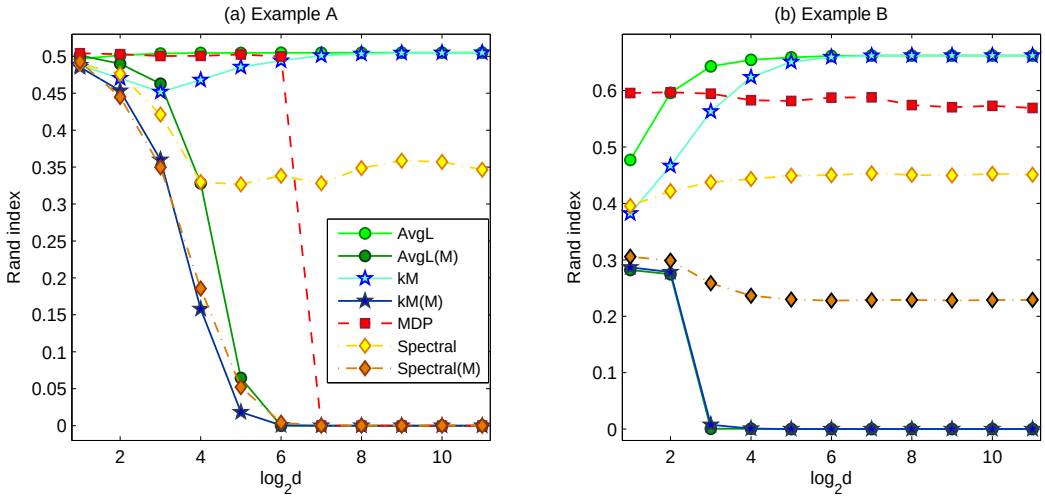

Figure 2: Rand indices for different algorithms in Examples A and B.

Failure of these popular algorithms shows the necessity to develop new methods for clustering high dimensional data. In this article, we use a new dissimilarity index called MADD (defined in Section 2) for this purpose, and develop cluster analysis based on it. It is interesting to note that both AvgL and kM yielded excellent results when, instead of Euclidean distance, MADD was used as the dissimilarity measure (see the curves corresponding to AvgL(M) and kM(M) in Figure 2). In both examples, they led to perfect clustering in high dimensions. It is well known that both AvgL and kM based on Euclidean distance are not much useful for finding non-convex clusters in the data, but Example B clearly shows that their MADD versions can overcome this limitation. Spectral clustering algorithm of Shi and Malik (2000) also performed well when a decreasing function of MADD (defined in Section 2) was used as the similarity measure. While the use of MADD led to perfect clustering for

large d in Example A, it significantly reduced the Rand index in Example B as well (see the curves corresponding to Spectral(M) in Figure 2).

The reasons behind the failure of Euclidean distance based clustering and the excellence of MADD based clustering are investigated in Section 2.We show that MADD uses the distance concentration property to its advantage. In this connection, we prove the high dimensional consistency of some popular clustering algorithms based on MADD. Simulation studies are also carried out to demonstrate superiority of these MADD based algorithms. In Section 3, we consider the problem of estimating the number of clusters from the data. This is an important problem in cluster analysis, and several methods are available for it. We will see that many of these methods perform better when MADD is used for their constructions. We investigate the high dimensional behavior of these MADD based estimation methods under appropriate regularity conditions. We also construct a new estimator based on penalized Dunn index and prove its high dimensional consistency. Empirical performance of different estimation methods are evaluated using simulation studies. Two benchmark data sets are analyzed in Section 4 for further comparison among different estimation methods and clustering algorithms. Finally, Section 5 gives a brief summary of the work and ends with some related discussions. All proofs and mathematical details are given in the Appendix.

2 Clustering algorithms based on MADD

Suppose that the whole sample $\mathcal{X} = \cup_{i=1}^{k_0} \mathcal{X}_i$ consists of n unlabeled observations from k_0 different populations, where $\mathcal{X}_i$ denotes the collection of n_i observations ($\sum_{i=1}^{k_0} n_i = n$) from the i-th ($i = 1, \ldots, k_0$) population. In Example A, we had $k_0 = 2$ and $n_1 = n_2 = 50$. In this example, for two observations $\mathbf{X} = (X^{(1)}, \ldots, X^{(d)})$ and $\mathbf{Y} = (Y^{(1)}, \ldots, Y^{(d)})$ from the second population, $\frac{1}{d}\|\mathbf{X} - \mathbf{Y}\|^2 = \frac{1}{d}\sum_{q=1}^{d}(X^{(q)} - Y^{(q)})^2$, being an average of an m-dependent sequence (with $m = 1$) of identically distributed random variables, converges to $E(X^{(1)} - Y^{(1)})^2 = 2\sigma_2^2 = 4$ in probability. But, if $\mathbf{X}$ comes from the first population and $\mathbf{Y}$ from the second population, $\frac{1}{d}\|\mathbf{X} - \mathbf{Y}\|^2$ converges in probability to $\sigma_1^2 + \sigma_2^2 + 1 = 3.5$. So,

for large d, all observations in $\mathcal{X}_2$ had their neighbors in $\mathcal{X}_1$, and clustering algorithms based on the Euclidean distance failed to put them in the same cluster. In Example B also, we have similar convergence of pairwise distances, which is given by the following lemma.

Lemma 1. *If* $\mathbf{X} \sim U_d(a_1, b_1)$, $\mathbf{Y} \sim U_d(a_2, b_2)$ *and they are independent, then* $\frac{1}{d}\|\mathbf{X} - \mathbf{Y}\|^2$ *converges in probability to* $b_1^2 + b_2^2$ *as the dimension* d *tends to infinity.*

From Lemma 1, it is clear that for large values of d, all observations in $\mathcal{X}_2$ and $\mathcal{X}_3$ had their neighbors in $\mathcal{X}_1$ only. So, AvgL and kM algorithms had miserable performance in this example as well.

In order to overcome this problem, we measure the dissimilarity between two observations $\mathbf{z}, \mathbf{w} \in \mathcal{X}$ by using **M**ean **A**bsolute **D**ifference of **D**istances (MADD), which is given by

$$\rho(\mathbf{z}, \mathbf{w}) = \frac{1}{n-2} \sum_{\mathbf{t} \in \mathcal{X} \setminus \{\mathbf{z}, \mathbf{w}\}} \big| \|\mathbf{z} - \mathbf{t}\| - \|\mathbf{w} - \mathbf{t}\| \big|. \tag{2}$$

In both Examples A and B, $d^{-1/2}\rho(\mathbf{z}, \mathbf{w}) \xrightarrow{P} 0$ as $d \to \infty$ if and only if $\mathbf{z}$ and $\mathbf{w}$ come from the same population. Otherwise, it converges to a positive constant. This leads to better performance by clustering algorithms based on MADD in high dimensions. MADD has some nice properties as a dissimilarity measure. Lemma 2 shows that it is a semi-metric.

Lemma 2. *For any collection* $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n\}$ *of* $n \geq 3$ *observations,* $\rho(\cdot, \cdot)$ *(as defined in (2)) is a semi-metric on* $\mathcal{X}$*.*

MADD is not a metric since it is possible to get $\mathbf{z} \neq \mathbf{w}$ such that $\rho(\mathbf{z}, \mathbf{w}) = 0$. But, when $\mathcal{X}$ consists of continuous random vectors, for $\mathbf{z}, \mathbf{w} \in \mathcal{X}$ and $\mathbf{z} \neq \mathbf{w}$, $\rho(\mathbf{z}, \mathbf{w}) > 0$ holds with probability one. Therefore, for all practical purposes, it behaves like a distance function. Using triangle inequality of the Euclidean metric, we get $\rho(\mathbf{z}, \mathbf{w}) \leq \|\mathbf{z} - \mathbf{w}\|$. So, closeness in terms of Euclidean distance also indicates closeness in terms of MADD, but not the converse. For high dimensional data, unlike Euclidean distance, MADD usually provides small dissimilarities among observations from the same population.

2.1 High dimensional behavior of MADD

To study the high dimensional behavior of ρ and associated clustering algorithms in details, we assume that $\mathcal{X} = \cup_{k=1}^{k_0} \mathcal{X}_k$ consists of n $(= \sum_{k=1}^{k_0} n_k)$ independent observations on the measurement vector $\mathbf{X} = (X^{(1)}, \ldots, X^{(d)})'$ coming from a mixture of k_0 populations, where $n_k = |\mathcal{X}_k| \geq 2$ for all $k = 1, 2, \ldots, k_0$, and the k-th population $(k = 1, \ldots, k_0)$ has mean vector $\boldsymbol{\mu}_{k,d}$ and scatter matrix $\boldsymbol{\Sigma}_{k,d}$. We also make the following assumptions.

$(A1)$ For each $k = 1, 2, \ldots, k_0$, fourth moments of the component variables $X^{(1)}, X^{(2)}, \ldots$ are uniformly bounded.

$(A2)$ There exists $\sigma_k^2 > 0$ for each k and $\nu_{kk'}$ for every $k \neq k'$ such that (a) $d^{-1} trace(\boldsymbol{\Sigma}_{k,d}) \to \sigma_k^2$ and (b) $d^{-1} \|\boldsymbol{\mu}_{k,d} - \boldsymbol{\mu}_{k,d'}\|^2 \to \nu_{kk'}$ as $d \to \infty$.

$(A3)$ For two independent observations $\mathbf{X}$ and $\mathbf{Y}$ from the k-th and the k'-th populations $(1 \leq k, k' \leq k_0)$, respectively, $\sum\sum_{1 \leq i < j \leq d} corr(Z^{(i)}, Z^{(j)})$ is of the order $\boldsymbol{o}(d^2)$, where $Z^{(i)} = (X^{(i)} - Y^{(i)})^2$ for $i = 1, 2, \ldots, d$.

Under $(A1) - (A3)$, high dimensional behavior of the Euclidean distance and MADD are given by the following lemma.

Lemma 3. *Suppose that we have a set of observations from k_0 populations satisfying $(A1) - (A3)$. If $\mathbf{X}$ and $\mathbf{Y}$ are two independent observations from the k-th and the k'-th populations respectively $(1 \leq k, k' \leq k_0)$, then*
(i) $d^{-1/2} \|\mathbf{X} - \mathbf{Y}\| \xrightarrow{P} (\sigma_k^2 + \sigma_{k'}^2 + \nu_{kk'})^{1/2}$ *as* $d \to \infty$ *(note that* $\nu_{kk'} = 0$ *for* $k = k'$*),*
(ii) $d^{-1/2} \rho(\mathbf{X}, \mathbf{Y}) \xrightarrow{P} c_{kk'}$ *as* $d \to \infty$*, where* $c_{kk'} \geq 0$*, and the equality holds if and only if* $\sigma_k^2 = \sigma_{k'}^2$ *and* $\nu_{jk} = \nu_{jk'}$ *for every* $j = 1, 2, \ldots k_0$ *(note that* $c_{kk'} = 0$ *for* $k = k'$*).*

The assumption $(A3)$ implies a form of weak dependence among the measurement variables. For time series data, $(A3)$ holds if the series has the ρ-mixing property, stationarity of the time series is not required. Hall et al. (2005) used a similar condition for studying high dimensional behavior of some popular classifiers. Recently, Biswas et al. (2014) assumed

this condition for investigating high dimensional consistency of their two-sample tests. The conditions $(A1)$ and $(A2)$ are quite common in the literature of high dimensional statistical analysis (see Hall et al. 2005), but $(A1)$ can be relaxed if instead of Euclidean distance, we use other distance functions (see Biswas et al. 2015). Note that the conditions $(A1)$ and $(A3)$ are only sufficient but not necessary to have weak law of large numbers (WLLN) for the sequence of dependent and non-identically distributed random variables $\{(X^{(i)}-Y^{(i)})^2; i \geq 1\}$. This convergence result may hold even when $(A3)$ is not satisfied, and we have seen that for Example B. The condition $(A2)$ helps us to compute the limiting value of $\frac{1}{d}\sum_{i=1}^{d}(X^{(i)}-Y^{(i)})^2$ as d diverges. If the component variables are independent and identically distributed (i.i.d), instead of $(A1)-(A3)$, we only need the second order moment of a component variable to be finite to have the results stated in Lemma 3.

From Lemma 3, it is clear that Euclidean distance can successfully differentiate two high dimensional populations only when their mean difference dominates the scale difference $(i.e., \nu_{kk'} > |\sigma_k^2 - \sigma_{k'}^2|)$. But, MADD can differentiate two populations whenever they differ in their means $(i.e., \nu_{kk'} > 0)$ and/or scatters $(i.e., \sigma_k^2 \neq \sigma_{k'}^2)$. Throughout this article, we will assume that the underlying population distributions differ in their locations and/or scales, and the groups of observations from each of these populations will be considered as the original clusters. On the other hand, groups obtained using clustering algorithms will be referred to as the estimated clusters.

2.2 High dimensional behavior of MADD based clustering

Lemma 3 helps us to explain the high dimensional behavior of different clustering algorithms observed in Section 1. Average linkage clustering algorithm (AvgL) begins with n groups, each consisting of a single observation. At each step, it chooses two closest groups and merges them into a single one. To measure closeness, the distance between two groups C_i and C_j is defined as $\Delta(C_i, C_j) = \frac{1}{|C_i||C_j|}\sum_{\mathbf{z}\in C_i, \mathbf{w}\in C_j}\|\mathbf{z}-\mathbf{w}\|$. AvgL stops merging when the pairwise distance between any two groups is bigger than a certain threshold or a specified number of

groups is achieved. The final groups thus formed are considered as the estimated clusters. Note that in Example A, for $\mathbf{X}_1, \mathbf{X}_2 \overset{i.i.d.}{\sim} N(\mathbf{0}_d, \sigma_1^2 \mathbf{\Sigma}_d)$ and $\mathbf{Y}_1, \mathbf{Y}_2 \overset{i.i.d.}{\sim} N(\boldsymbol{\mu}_d, \sigma_2^2 \mathbf{\Sigma}_d)$, we have $P(\|\mathbf{X}_1 - \mathbf{X}_2\| < \|\mathbf{X}_1 - \mathbf{Y}_1\| < \|\mathbf{Y}_1 - \mathbf{Y}_2\|) \to 1$ as $d \to \infty$ (follows from Lemma 3). So, for high dimensional data, in the first 49 steps, all observations from the first population were merged into a single group, and in each subsequent steps, one observation from the second population was added to it. As a result, when AvgL ended with two estimated clusters, one of them had a single observation from the second population, and the other had the rest of the observations. This led to a Rand index of 0.505. A similar phenomenon occurred in Example B as well, where two of the three clusters estimated by AvgL had one observation each from the third population (i.e., $U_d(2, 2.5)$), while the third contained the rest. As a result, the Rand index turned out to be 0.662. We observed the same phenomenon even when single or complete linkage was used instead of AvgL.

We observed a diametrically opposite behavior for AvgL(M), where MADD is used in place of Euclidean distance to define the distance between two groups. Note that in Example A, as $d \to \infty$, after appropriate scaling, both $\rho(\mathbf{X}_1, \mathbf{X}_2)$ and $\rho(\mathbf{Y}_1, \mathbf{Y}_2)$ converge to 0, while $\rho(\mathbf{X}_1, \mathbf{Y}_1)$ converges to a positive constant. Therefore, any linkage method eventually leads to two estimated clusters which completely separate the two populations. Because of this perfect clustering, AvgL(M) had zero Rand index for higher values of d. Similar convergence occurs in Example B as well. The following theorem shows that if the underlying populations satisfy $(A1)-(A3)$, AvgL(M) will always lead to perfect clustering for high dimensional data.

Theorem 1. *Suppose that we have a set of observations from k_0 different populations satisfying $(A1) - (A3)$. Also assume that $\sigma_k^2 \neq \sigma_{k'}^2$ and/or $\nu_{kk'} > 0$ for every $1 \leq k \neq k' \leq k_0$. If AvgL(M) is used to estimate k_0 clusters in the data, its Rand index will converge to 0 in probability as d tends to infinity.*

Now, let us consider the case of k-Means clustering algorithm (kM). It divides the whole sample into disjoint groups in such a way that an observation is closer to its own group center than the center of any other group. For a given k, it starts with an initial choice of k groups

and then at each step, an observation $\mathbf{x}$ is assigned to the group having the center closest to $\mathbf{x}$. Group centers are updated accordingly. This iterative process is terminated when no groups are modified further. Mathematically, kM algorithm aims at finding k groups $C_1, \ldots, C_k$ with centers $\mathbf{m}_1, \ldots, \mathbf{m}_k$ such that $\varphi(C_1, \ldots, C_k) = \sum_{j=1}^{k} \sum_{i:\mathbf{x}_i \in C_j} \|\mathbf{x}_i - \mathbf{m}_j\|^2$ is minimized. Now, using the distance convergence results in Lemma 3, one can show that in Examples A and B, for large d, φ will be minimized when we have the same type of estimated clusters as obtained in the case of AvgL. So, in high dimensions, AvgL and kM had similar behavior. Note that $\varphi(C_1, \ldots, C_k)$ can be re-written as $\varphi(C_1, \ldots, C_k) = \sum_{r=1}^{k} \frac{1}{2|C_r|} \sum_{\mathbf{z},\mathbf{w} \in C_r} \|\mathbf{z} - \mathbf{w}\|^2$. Using this formulation, we construct the MADD version of kM (denoted by kM(M)), where we minimize $\varphi^*(C_1, \ldots, C_k) = \sum_{r=1}^{k} \frac{1}{2|C_r|} \sum_{\mathbf{z},\mathbf{w} \in C_r} \rho^2(\mathbf{z}, \mathbf{w})$. For this minimization, again we start with k initial groups and use an iterative algorithm. At each step, distance of an observation $\mathbf{x}$ from a group C_j is computed as $\rho_0(\mathbf{x}, C_j) = \frac{1}{|C_j|} \sum_{\mathbf{z} \in C_j} \rho^2(\mathbf{x}, \mathbf{z})$. The observation is assigned to the group $C_{\tilde{k}}$, where $\tilde{k} = \operatorname{argmin} \rho_0(\mathbf{x}, C_k)$. This is done for all observations, and the process is repeated until convergence. From Lemma 3, it is clear that for k_0 estimated clusters $C_1, \ldots, C_{k_0}$, $\varphi^*(C_1, \ldots, C_{k_0})$ attains its minimum value zero if and only if we have perfect clustering. So, as expected, kM(M) had excellent performance in Examples A and B, particularly when d was large. This perfect clustering phenomenon of kM(M) is asserted by the following theorem.

Theorem 2. *Suppose that we have a set of observations from k_0 different populations satisfying $(A1) - (A3)$. Also assume that $\sigma_k^2 \neq \sigma_{k'}^2$ and/or $\nu_{kk'} > 0$ for every $1 \leq k \neq k' \leq k_0$. If kM(M) is used to estimate k_0 clusters in the data, its Rand index will converge to 0 in probability as d tends to infinity.*

Remark 1. *We only need the distance convergence results in Lemma 3 to prove Theorems 1 and 2. Recall that the conditions $(A1) - (A3)$ are only sufficient but not necessary for proving the convergence results in Lemma 3. So, Theorems 1 and 2 may continue to hold even when $(A1) - (A3)$ are not satisfied.*

The spectral clustering algorithm of Shi and Malik (2000) also failed to perform well in the examples considered in Section 1. Note that spectral clustering methods deal with an edge-

weighted graph with nodes $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n\}$ and a symmetric weight matrix $\mathcal{W} = ((w_{ij}))_{d\times d}$, where w_{ij} represents similarity between $\mathbf{x}_i$ and $\mathbf{x}_j$. The matrix $\mathcal{W}$ is usually computed from a similarity matrix $\mathcal{S}$, and different methods are available for it (see von Luxburg 2007). Often $\mathcal{S} = ((s_{ij}))$ itself is used as $\mathcal{W}$, and one popular choice for s_{ij} is the radial basis function $s_{ij} = \exp\{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}\}$, where σ is a tuning parameter that controls the degree of similarity. These algorithms implicitly assume that s_{ij} will be large (respectively, small) if $\mathbf{x}_i$ and $\mathbf{x}_j$ belong to the same population (respectively, different populations). Unfortunately that was not the case in Examples A and B. Because of this violation of cluster assumption, spectral clustering algorithm had poor performance. However, this assumption holds if s_{ij} is defined using $\rho^2(\mathbf{x}_i, \mathbf{x}_j)$ instead of $\|\mathbf{x}_i - \mathbf{x}_j\|^2$. So, Spectral(M), the spectral clustering algorithm based on MADD, had much improved performance in these examples, especially for higher values of d.

The clustering algorithm based on MDP distance (Ahn et al. 2012) largely depends on the data piling property, which occurs only when the dimension exceeds the sample size. So, as expected, it performed poorly in both examples for smaller values of d. Surprisingly, it failed in Example B even when d was large. A simple investigation explains this artifact. MDP clustering algorithm estimates the clusters by using binary splits at each step. For observations from two populations, for which the distance convergence results in Lemma 3 hold, this algorithm consistently separates the two populations in HDLSS set up when

$$\nu_{12} + \frac{\sigma_1^2}{n_1} + \frac{\sigma_2^2}{n_2} > \min\left\{\frac{G+n_1}{Gn_1}\sigma_1^2, \frac{G+n_2}{Gn_2}\sigma_2^2\right\}, \tag{3}$$

where G is a pre-specified minimum number of observations in a cluster, and σ_1^2, σ_2^2 and ν_{12} have the same meaning as in $(A2)$. Following Ahn et al. (2012), we used $G = 5$ throughout this article. Recall that in Example B, all three populations had the same location, and the condition (3) was violated for each pair of classes. So, MDP clustering could not perform well in that example.

2.3 Comparison of clustering algorithms using simulated datasets

Using Examples A and B, we have already demonstrated the importance of MADD for clustering high dimensional data. Here, we analyze six simulated data sets for further evaluation of MADD based clustering algorithms. In each case, we generated the data sets taking 50 observations from each population, and different clustering algorithms were used on these data sets assuming the number of clusters to be known. For each example, we considered $d = 100, 200$ and 500, and each experiment was repeated 100 times. Average Rand indices of different algorithms are computed over these 100 trials, and they are reported in Table 1. MDP clustering method needs the number of eigen-vectors T to be specified. We used $T = 1, 2, 3$ as in Ahn et al. (2012), and in each case, the best result is reported. Descriptions of the six data sets are given below.

Example-1: Observations were generated from three Gaussian distributions with the same scatter matrix $\mathbf{\Sigma}_d^\circ = ((0.5^{|i-j|}))_{d\times d}$ but different mean vectors $\boldsymbol{\mu}_1$, $\boldsymbol{\mu}_2$ and $\boldsymbol{\mu}_3$, respectively. We took $\boldsymbol{\mu}_1 = \mathbf{0}_d$, while $\boldsymbol{\mu}_2$ (respectively, $\boldsymbol{\mu}_3$) had the first $d/2$ elements equal to 0.75 (respectively, -0.75) and the rest equal to 0.

Example-2: We used observations from four Gaussian distributions $N(\boldsymbol{\alpha}, \mathbf{\Sigma}_d^\circ)$, $N(\boldsymbol{\beta}, 4\mathbf{\Sigma}_d^\circ)$, $N(-\boldsymbol{\alpha}, \mathbf{\Sigma}_d^\circ)$ and $N(-\boldsymbol{\beta}, 4\mathbf{\Sigma}_d^\circ)$, which differed in their locations and scales. The mean vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)'$ had $\alpha_i = 1$ and $\alpha_i = 0.5$ for even and odd values of i, respectively. We took $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_d)'$ with $\beta_i = (-1)^i \alpha_i$ for $i = 1, 2, \ldots, d$, and $\mathbf{\Sigma}_d^\circ$ as in Example-1.

Example-3: We considered three uniform distributions with disjoint supports S_1, S_2 and S_3, where S_i $(i = 1, 2, 3)$ is given by $S_i = \{\mathbf{x} \in \mathbb{R}^d : \ i - 1 \le \mathbf{x}'\mathbf{\Sigma}_d^{\circ -1}\mathbf{x} \le i - 1/2\}$, and $\mathbf{\Sigma}_d^\circ$ is as in Example-1. Figure 3(a) shows the supports of the three distributions for $d = 2$.

Example-4: Define three non-convex sets $S_1^\circ = \{(x, y) : y \ge 0, 1 \le \sqrt{(x-2)^2 + y^2} \le 1.5\}$, $S_2^\circ = \{(x, y) : y \ge 0, 1 \le \sqrt{(x+2)^2 + y^2} \le 1.5\}$ and $S_3^\circ = \{(x, y) : y \le 0, 4 \le \sqrt{x^2 + y^2} \le 4.5\}$ (see Figure 3(b)). We generated $d/2$ independent bivariate observations from the uniform distribution on S_i° $(i = 1, 2, 3)$ to get d components of an observation from the i-th population.

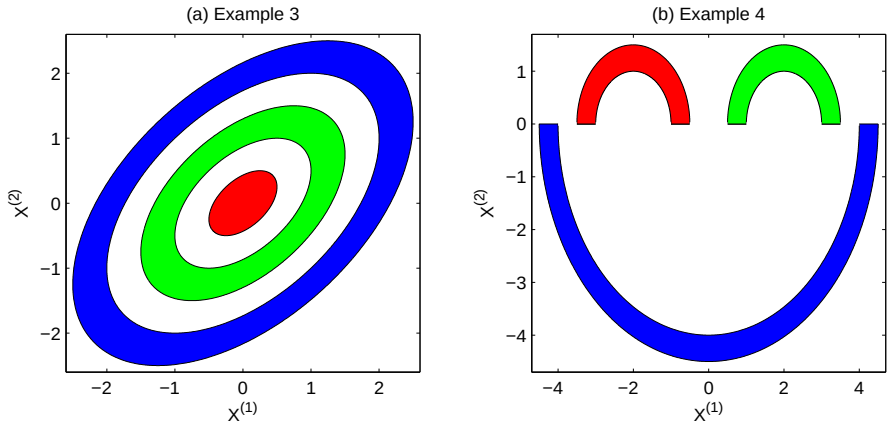

Figure 3: Different populations in Examples 3 and 4 when $d = 2$.

Example-5: Observations were generated from two auto-regressive processes $X_t = 0.75 + 0.25X_{t-1} + \varepsilon_t$ and $X_t = 0.25 + 0.75X_{t-1} + \varepsilon_t$ for $t = 1, 2, \ldots, d$. In both cases, we had $\varepsilon_t \sim N(0, 1)$ for every t. The distribution of X_0 was taken to be $N(1, 16/15)$ and $N(1, 16/7)$ in these two cases to make the processes stationary.

Example-6: Let $\mathcal{S}_d$ be the d-dimensional unit sphere with center at the origin, and $\mathcal{C}_d$ be the largest hypercube inscribed in it. We considered two uniform distributions, one on $\mathcal{S}_d$ and the other on $\mathcal{C}_d$. Note that if $\mathbf{X}$ comes from the first population, then $P(\mathbf{X} \in \mathcal{C}_d) \to 0$ as $d \to \infty$. So, the two populations become separated in high dimensions.

Table 1 clearly shows that both average linkage and k-means clustering performed better when MADD was used. AvgL(M) and kM(M) had the best performance in almost all cases, and they led to perfect clustering for higher values of d in all six data sets. MDP clustering algorithm performed poorly for $d = 100$. For $d = 200$ and 500, it performed well in Examples 1, 2 and 4, but in the other three examples, where the population distributions did not differ in their means, it had miserable performance. For spectral clustering also, we got better results when MADD was used. In Example-6, when the clustering algorithms based on Euclidean distance and MDP distance had Rand indices close to 0.5, those based on MADD led to perfect clustering for all values of d considered here. Among the algorithms based on MADD, overall performance of AvgL(M) and kM(M) was much better than Spectral(M).

Table 1: Average Rand indices of different clustering algorithms on simulated data sets

	d	AvgL	AvgL(M)	kM	kM(M)	MDP	Spectral	Spectral(M)
Ex-1	100	0.2906	0.0865	**0.0185**	0.0367	0.5801	0.2512	0.1851
	200	0.2168	0.0104	0.0201	0.0095	**0.0000**	0.2419	0.1953
	500	0.0429	**0.0000**	0.0074	**0.0000**	**0.0000**	0.2330	0.1919
Ex-2	100	0.7335	0.0502	0.4206	**0.0067**	0.6204	0.2609	0.0415
	200	0.7361	0.0115	0.6206	**0.0001**	0.0714	0.2462	0.0440
	500	0.7378	**0.0000**	0.6982	**0.0000**	0.0018	0.2187	0.0434
Ex-3	100	0.6616	**0.0000**	0.6608	**0.0000**	0.5885	0.4492	0.2348
	200	0.6619	**0.0000**	0.6617	**0.0000**	0.5826	0.4513	0.2298
	500	0.6619	**0.0000**	0.6619	**0.0000**	0.5761	0.4515	0.2286
Ex-4	100	0.2346	**0.0000**	0.2762	**0.0000**	0.5841	0.0745	0.0417
	200	0.2330	**0.0000**	0.2769	**0.0000**	**0.0000**	0.0714	0.0361
	500	0.2327	**0.0000**	0.2748	**0.0000**	**0.0000**	0.0752	0.0478
Ex-5	100	0.5047	0.3516	0.5027	**0.2271**	0.4895	0.4801	0.2584
	200	0.5048	**0.0762**	0.5040	0.0784	0.4863	0.4813	0.1231
	500	0.5048	**0.0028**	0.5048	0.0060	0.4795	0.4726	0.0119
Ex-6	100	0.5047	**0.0000**	0.5042	**0.0000**	0.4857	0.5000	**0.0000**
	200	0.5048	**0.0000**	0.5048	**0.0000**	0.4836	0.4992	**0.0000**
	500	0.5048	**0.0000**	0.5048	**0.0000**	0.4803	0.4992	**0.0000**

Bold figures indicate best result in each example.

3 Estimation of number of clusters

So far we have assumed the number of clusters k_0 to be known for our analysis. But in practice, one needs to estimate k_0 from the data. This is a challenging problem in cluster analysis, and several estimation methods have been proposed for it (see Calinski and Harabasz 1974; Hartigan 1975; Krzanowski and Lai 1985; Kaufman and Rousseeuw 1990; Tibshirani et al. 2001; Sugar and James 2003; Wang 2010). Brief descriptions of some of these methods that we use in this article are given below. These estimation methods can be used with any base clustering algorithm. Throughout this article, we use average linkage or k-means algorithm (either based on Euclidean distance or based on MADD) for base clustering.

KL statistic (Krzanowski and Lai 1985): For a given k, if $C_1, \ldots, C_k$ are the k clusters estimated by the base clustering algorithm, then KL statistic is given by $\mathrm{KL}(k) = \left|\frac{\mathrm{DIFF}(k)}{\mathrm{DIFF}(k+1)}\right|$, where $\mathrm{DIFF}(k) = (k-1)^{2/d}W_{k-1} - k^{2/d}W_k$, and $W_k = \sum_{j=1}^{k} \frac{1}{2|C_j|} \sum_{\mathbf{z},\mathbf{w}\in C_j} \|\mathbf{z}-\mathbf{w}\|^2$ is the

within group sum of squares. $\mathrm{KL}(k)$ is computed for a range of values $\{2, 3, \ldots, K\}$ of k and $\hat{k}_{KL} = \mathrm{argmax}_{2 \le k \le K} \mathrm{KL}(k)$ is used to estimate k_0.

Gap statistic (Tibshirani et al. 2001): For any fixed k, the Gap statistic is defined as $\mathrm{Gap}(k) = \frac{1}{B}\sum_{b=1}^{B} \log(W_k^{(b)}) - \log(W_k)$, where W_k has the same meaning as in KL statistic, and $W_k^{(b)}$ is the within group sum of squares computed using the b-th bootstrap sample $(b = 1, \ldots, B)$ generated from a reference distribution. The number of clusters k_0 is estimated by $\hat{k}_G = \min\{k : \ \mathrm{Gap}(k) \ge \mathrm{Gap}(k+1) - s_{k+1}\}$, where $s_k = \sqrt{(1 + \frac{1}{B})} sd_k$, and sd_k is the standard deviation of $W_k^{(b)}$. Unlike KL statistic, $\mathrm{Gap}(k)$ can be defined for $k = 1$ as well.

Jump statistic (Sugar and James 2003): For $k \ge 1$, the Jump statistic is defined as $\mathrm{Jump}(k) = \hat{d}_k^{-t} - \hat{d}_{k-1}^{-t}$, where $\hat{d}_0^{-t} = 0$, $\hat{d}_k = \frac{1}{d}\sum_{i=1}^{n} \min_{r=1,2,\ldots,k} (\mathbf{x}_i - \mathbf{m}_r)' \Gamma^{-1} (\mathbf{x}_i - \mathbf{m}_r)$ for $k \ge 1$, and $\mathbf{m}_r$ is the center of the r-th cluster. The authors suggested to use $\Gamma = \boldsymbol{I}_d$ (the $d \times d$ identity matrix) and $t = \frac{d}{2}$. The number of clusters is estimated by $\hat{k}_J = \mathrm{argmax}_{1 \le k \le K} \mathrm{Jump}(k)$. Note that for k-means clustering with $\Gamma = \boldsymbol{I}_d$, we get $\hat{d}_k = \frac{1}{d}\sum_{i=1}^{n} \min_{r=1,2,\ldots,k} \|\mathbf{x}_i - \mathbf{m}_r\|^2 = \frac{1}{d}\sum_{j=1}^{k}\sum_{\mathbf{z} \in C_j} \|\mathbf{z} - \mathbf{m}_j\|^2 = \frac{1}{d}\sum_{j=1}^{k} \frac{1}{2|C_j|}\sum_{\mathbf{z},\mathbf{w} \in C_j} \|\mathbf{z} - \mathbf{w}\|^2 = W_k/d$.

Cross-validated Rand index (Wang 2010): The whole sample $\mathcal{X}$ is divided into three parts $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}$ and $\mathcal{X}^{(3)}$ of sizes m, m and $n - 2m$, respectively. For any given k, the first two parts are used to develop two clustering algorithms $\psi_1 = \delta_{\mathcal{X}^{(1)},k}(\cdot)$ and $\psi_2 = \delta_{\mathcal{X}^{(2)},k}(\cdot)$, which are then used on the observations in $\mathcal{X}^{(3)}$ to estimate clustering instability given by $\mathrm{Ins}(k) = \binom{n-2m}{2}^{-1} \sum_{\mathbf{x} \ne \mathbf{y} \in \mathcal{X}^{(3)}} \mathbb{I}\Big[\mathbb{I}\{\psi_1(\mathbf{x}) = \psi_1(\mathbf{y})\} + \mathbb{I}\{\psi_2(\mathbf{x}) = \psi_2(\mathbf{y})\} = 1\Big]$. This process is repeated B times. The author proposed two methods to aggregate these results. In one method, the average instability over B repetitions is computed, and k_0 is estimated by minimizing this average instability with respect to k. This method is called cross-validation using average (CV_a). In the other method, for each repetition the number of clusters is estimated by minimizing $\mathrm{Ins}(k)$ over k, and finally the modal value of the minimizer is used as an estimator of k_0. This method is called cross-validation using vote (CV_v).

Along with these methods, we use another method based on Dunn index (Dunn 1973). For a fixed k, let $C_1, \ldots, C_k$ be the clusters estimated by a base clustering algorithm. Define

$\Delta_0(C_i) = \frac{1}{|C_i|(|C_i|-1)} \sum_{\mathbf{z},\mathbf{w}\in C_i} \|\mathbf{z}-\mathbf{w}\|$ and $\Delta(C_i, C_j) = \frac{1}{|C_i||C_j|} \sum_{\mathbf{z}\in C_i,\mathbf{w}\in C_j} \|\mathbf{z}-\mathbf{w}\|$ (instead of average, other suitable measures can also be used, see Dunn 1973). Then the Dunn index is given by $\mathrm{DI}(k) = B_k^\circ / W_k^\circ$, where $W_k^\circ = \max_{1\le i\le k} \Delta_0(C_i)$ and $B_k^\circ = \min_{1\le i<j\le k} \Delta(C_i, C_j)$, respectively. Dunn (1973) used this index for cluster validation, but here we use it to estimate the number of clusters by $\hat{k}_D = \mathrm{argmax}_{2\le k\le K}\, DI(k)$.

When AvgL or kM is used for base clustering, we use the usual versions of these statistics based on Euclidean norm to estimate k_0. However, when AvgL(M) or kM(M) is used for base clustering, we use their MADD versions, where the Euclidean norm is replaced by MADD. For instance, ρ is used in place of $\|\cdot\|$ to define W_k for KL, Gap and Jump statistics, and to define Δ_0 and Δ for Dunn index.

To investigate the performance of these estimation methods, we considered the six examples used in Section 2.3 with $d = 500$. For the methods based on cross-validated Rand index, we used $B = 100$, and m was taken to be the largest multiple of 5 not exceeding $\frac{n}{3}$. For the Gap statistic, we used $B = 100$, and bootstrap samples were generated from the uniform distribution on the range of the measurement variables. When d was large (in the order of hundreds or more), use of $t = \frac{d}{2}$ led to poor performance by the method based on Jump statistic. So, we tried several choices of t, and based on our empirical experience, we selected $t = 1$ when MADD was used. However, we were unable to find any such t when the Euclidean norm was used. In such cases, we performed our experiments with several choices of t and reported the best results. Throughout this article, we considered the values of k in the range $\{1, 2, \ldots, 12\}$. However, only Gap and Jump statistics are defined for $k = 1$.

Figures 4 and 5 show barplots for the number of clusters estimated by different methods in Examples 1 and 2. From these figures, it is clear that barring the method based on Gap statistic, all other methods worked better when MADD was used. We observed similar phenomenon in other examples as well. Therefore, we have only reported the results for the MADD versions of these methods in Table 2. For the method based on Gap statistic, there was no clear winner. Gap statistic based on Euclidean distance (respectively, MADD)

showed a tendency to underestimate (respectively, overestimate) the number of clusters in some cases. For this method, we used both the Euclidean distance and MADD in all examples, and in each case, the best result has been reported in Table 2. This table clearly shows that except for Examples 1 and 6, Gap statistic yielded very poor performance throughout. CV_a and CV_v also underestimated the number of clusters in Examples 2 and 3. But the other three methods based on KL statistic, Jump statistic and Dunn index correctly estimated the number of clusters on all occasions.

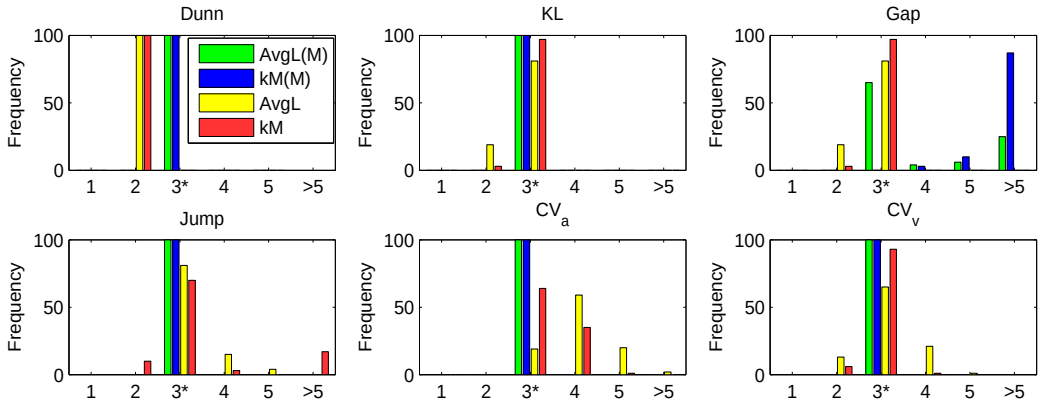

Figure 4: Barplots for number of clusters estimated by different methods in Example 1

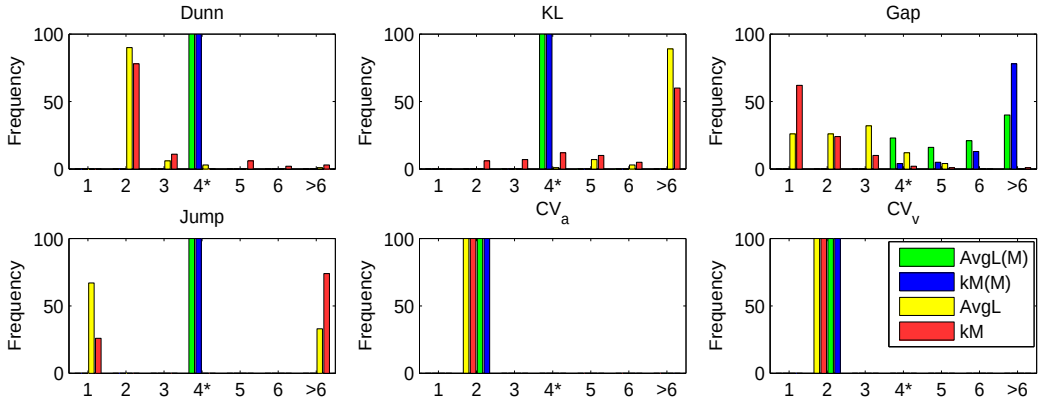

Figure 5: Barplots for number of clusters estimated by different methods in Example 2

The success of these three methods motivated us to carry out a theoretical investigation on their high dimensional behavior. For this investigation, we consider $(A1)$ and $(A2)$ as before. But, instead of $(A3)$, we assume a slightly stronger condition, which is given below.

Table 2: Frequency distribution for the estimated number of clusters

		Average linkage												k-means										
	k	1	2	3	4	5	6	7	8	9	10	11	12	1	2	3	4	5	6	7	8	9	10	11	12
Ex-1	Dunn*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	KL*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	Gap?	0	19	**81**	0	0	0	0	0	0	0	0	0	0	3	**97**	0	0	0	0	0	0	0	0	0
	Jump*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	CV_a^*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	CV_v^*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
Ex-2	Dunn*	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0
	KL*	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0
	Gap?	0	0	0	**23**	16	21	15	12	5	1	5	2	0	0	0	**4**	5	13	20	14	18	10	10	6
	Jump*	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0
	CV_a^*	0	100	0	**0**	0	0	0	0	0	0	0	0	0	100	0	**0**	0	0	0	0	0	0	0	0
	CV_v^*	0	100	0	**0**	0	0	0	0	0	0	0	0	0	100	0	**0**	0	0	0	0	0	0	0	0
Ex-3	Dunn*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	KL*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	Gap?	0	0	**12**	13	41	18	13	3	0	0	0	0	0	0	**0**	2	20	35	23	18	1	1	0	0
	Jump*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	CV_a^*	0	100	**0**	0	0	0	0	0	0	0	0	0	0	100	**0**	0	0	0	0	0	0	0	0	0
	CV_v^*	0	100	**0**	0	0	0	0	0	0	0	0	0	0	100	**0**	0	0	0	0	0	0	0	0	0
Ex-4	Dunn*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	KL*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	Gap?	0	20	**19**	10	15	9	9	4	9	1	2	2	0	3	**27**	23	24	9	4	6	2	0	1	1
	Jump*	0	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0
	CV_a^*	0	9	**91**	0	0	0	0	0	0	0	0	0	0	9	**91**	0	0	0	0	0	0	0	0	0
	CV_v^*	0	9	**91**	0	0	0	0	0	0	0	0	0	0	9	**91**	0	0	0	0	0	0	0	0	0
Ex-5	Dunn*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	KL*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	Gap?	0	**19**	12	29	28	10	2	0	0	0	0	0	0	**0**	1	17	47	28	7	0	0	0	0	0
	Jump*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	CV_a^*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	CV_v^*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
Ex-6	Dunn*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	KL*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	Gap?	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	Jump*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	CV_a^*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0
	CV_v^*	0	**100**	0	0	0	0	0	0	0	0	0	0	0	**100**	0	0	0	0	0	0	0	0	0	0

Figures in bold indicates the frequency corresponding to the actual number of clusters

* Results based on MADD versions

? Both Euclidean and MADD versions were used and the best result is reported

($A3^\circ$) Let $\mathbf{X}$ and $\mathbf{Y}$ be two independent observations from the k-th and the k'-th populations ($1 \leq k, k' \leq k_0$), respectively. Then, $\sum\sum_{1\leq i<j\leq d} corr(Z^{(i)}, Z^{(j)})$ is of the order $\boldsymbol{O}(d^{2(1-\varepsilon)})$ for some $\varepsilon > 0$, where $Z^{(i)} = (X^{(i)} - Y^{(i)})^2$ $(i = 1, 2, \ldots, d)$.

Under these regularity conditions, the three methods have some nice properties in high dimensions if an appropriate base clustering algorithm is used. To make sense of what we mean by an appropriate base clustering algorithm, we now introduce the concept of perfect and order preserving (POP) clustering.

Definition 1. *For any fixed k, let $C_1^{(k)}, C_2^{(k)}, \ldots, C_k^{(k)}$ be k clusters estimated using a clustering algorithm δ on $\mathcal{X} = \cup_{i=1}^{k_0}\mathcal{X}_i$, which consists of observations from k_0 classes. δ is called perfect and order preserving (POP) at k_0 if the following conditions hold.*

(*a*) *For $k = k_0$, the clustering algorithm is perfect (i.e., $C_i^{(k_0)} = \mathcal{X}_{\pi(i)}$ for all $i = 1, 2, \ldots, k_0$ and some permutation π of $\{1, 2, \ldots, k_0\}$).*

(*b*) *For any $k < k_0$ and for every $i = 1, 2, \ldots, k_0$, there exists $j \leq k$ such that $C_i^{(k_0)} \subseteq C_j^{(k)}$.*

(*c*) *For any $k > k_0$ and for every $i = 1, 2, \ldots, k$, there exits $j \leq k_0$ such that $C_i^{(k)} \subseteq C_j^{(k_0)}$.*

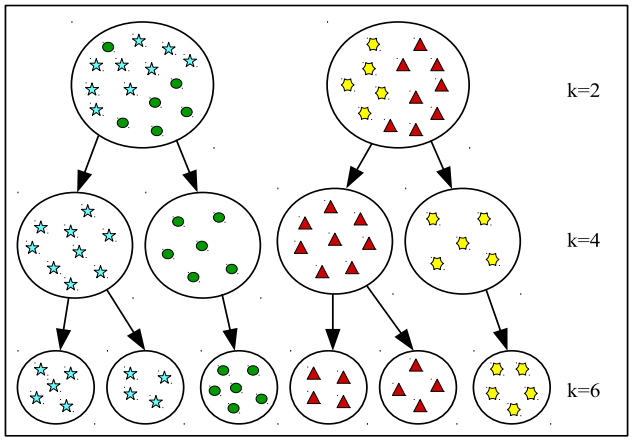

Figure 6: A clustering algorithm which is POP at 4.

Figure 6 demonstrates a POP clustering (at 4) by taking only one value of k smaller than 4 and one value of k bigger than 4. But, one should notice that property (b) (respectively,

property (c)) has to hold for all $k < 4$ (respectively, $k > 4$). It is easy to check that any hierarchical algorithm is order preserving (i.e., satisfies (b) and (c)). So, if it leads to perfect clustering at k_0, it becomes POP at k_0. In Theorems 1 and 2, we have seen that AvgL(M) and kM(M) become perfect with probability tending to one as d tends to infinity. Using Lemma 3(ii), one can show that, in fact, they become POP at k_0 with probability tending to one as the dimension diverges. Assuming that such a POP algorithm is used for base clustering, the following theorem shows the high dimensional behavior of estimators based on the MADD versions of Dunn index, KL statistic and JUMP statistic.

Theorem 3. *Suppose that there are observations from $k_0 \geq 2$ different populations, which satisfy $(A1), (A2)$ and $(A3^\circ)$, and also assume that the base clustering algorithm is POP at k_0. Then (i) $\hat{k}_D^* \xrightarrow{P} k_0$, (ii) $\hat{k}_{KL}^* \xrightarrow{P} k_0$ and (iii) $P(\hat{k}_J^* \geq k_0) \to 1$ as $d \to \infty$.*
(Here $\hat{k}_D^*$, $\hat{k}_{KL}^*$ and $\hat{k}_J^*$ are the number of clusters estimated by MADD versions of Dunn index, KL *statistic and* JUMP *statistic (with $t = 1$), respectively.)*

Theorem 3 shows the high dimensional consistency of the methods based on KL statistic and Dunn index. But, since DI(1) and KL(1) are not defined, none of these methods can detect the presence of a single cluster. JUMP statistic can be used in such situations. But, Theorem 3 only shows that $P(\hat{k}_J^* \geq k_0) \to 1$ as $d \to \infty$. So, it can overestimate the number of clusters in some cases. In order to overcome these limitations of MADD based procedures, we define a penalized version of Dunn index (PDI). For any fixed k, it is given by $\text{PDI}_\lambda(k) = B_k^\circ / W_k^\circ - \lambda k \log(d)$, where W_k° and B_k° (for $k > 1$) have the same meaning as in Dunn index, $B_1^\circ \stackrel{def}{=} B_2^\circ$ and $\lambda > 0$ is a parameter that controls the penalty. We estimate the number of clusters by maximizing $\text{PDI}_\lambda(k)$ with respect to k and denote it by $\hat{k}_{PD}$ ($\hat{k}_{PD}^*$ when its MADD version is used). The following theorem shows the high dimensional consistency of $\hat{k}_{PD}^*$ for any fixed choice of λ.

Theorem 4. *Suppose that there are observations from $k_0 \geq 1$ population(s), which satisfy $(A1), (A2)$ and $(A3^\circ)$. If the base clustering algorithm is POP at k_0, then $\hat{k}_{PD}^* \xrightarrow{P} k_0$ as $d \to \infty$.*

Remark 2. *In Theorems 3 and 4, the assumption that the base clustering algorithm is POP at k_0 only helps us to derive asymptotic orders of W_k, B_k, W_k° and B_k° for different values of k. These orders are maintained if AvgL(M) or kM(M) is used for base clustering, So, $\hat{k}_D^*, \hat{k}_{KL}^*$ and $\hat{k}_{PD}^*$ have high dimensional consistency in those situations as well.*

To evaluate the empirical performance of the estimator based on the MADD version of PDI, we used it (with $\lambda = 0.015$) on Examples $1 - 6$. In all these cases, it led to the same result as obtained using the MADD version of Dunn index. For further evaluation of its performance, we generated 100 observations from a uniform distribution on a unit hypercube of 500 dimension, and repeated the experiment 100 times. In all these 100 cases, this method successfully identified the presence of a single cluster in the data set. In this example, the JUMP statistic based on MADD and the GAP statistic based on Euclidean norm also estimated k_0 correctly on all occasions.

4 Analysis of benchmark data sets

We analyzed two benchmark data sets, 'Lymphoma' data and 'Synthetic Control Chart' data, for further evaluation of our proposed methods. Lymphoma data set was first analyzed by Alizadeh et al. (2000) for identification of distinct types of lymphoma, and it is available in the `R` package `spls`. Control Chart data set can be obtained from the `UCI Machine Learning Repository` (https://archive.ics.uci.edu/ml/datasets.html).

4.1 Lymphoma data

This data set contains expression levels of 4026 genes for 42 diffuse large B-cell lymphoma (DLBCL), 9 follicular lymphoma (FL) and 11 chronic lymphocytic leukemia (CLL) samples. A plot of these 62 observations is given in Figure 7.

We used different methods to estimate the number of clusters in this data set and the results are given in Table 3. From this table we can see that all methods, except for those

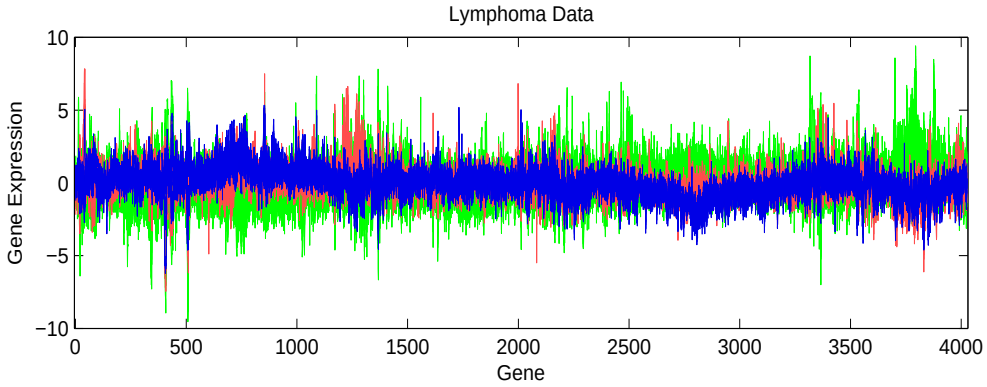

Figure 7: Gene expression of 4026 genes for DLBCL(■), FL(■) and CLL(■).

based on Gap statistic, identified two clusters in this data set. When we used different clustering algorithms to find two clusters in the data, all of them put almost all observations from the class DLBCL in one cluster and the observations from FL and CLL classes in the other cluster (see Figure 8(a)). This gives an indication that FL and CLL classes are very hard to differentiate from unlabeled data, which can be seen in Figure 7 as well. This claim can also be justified by the behavior of FL, which can sometimes present itself as CLL (see https://en.wikipedia.org/wiki/B-cell_chronic_lymphocytic_leukemia).

Table 3: Number of clusters estimated by different methods in 'Lymphoma' data

	Dunn	PDI	KL	Gap	Jump	CV_a	CV_v
AvgL	2	2	2	12	2	2	2
AvgL(M)	2	2	2	7	2	2	2
kM	2	2	2	12	2	2	2
kM(M)	2	2	2	7	2	2	2

Since it was known that the observations were actually from three populations, we used different clustering algorithms to find three clusters in this data set as well. In Figure 8(b), one can see that both AvgL and kM failed to identify the three populations. They could only form a new cluster consisting of a single observation from the class DLBCL. But, MADD versions of these algorithms successfully differentiated between observations from FL and CLL classes. The method based on MDP led to perfect clustering, but spectral clustering algorithms did not perform well.

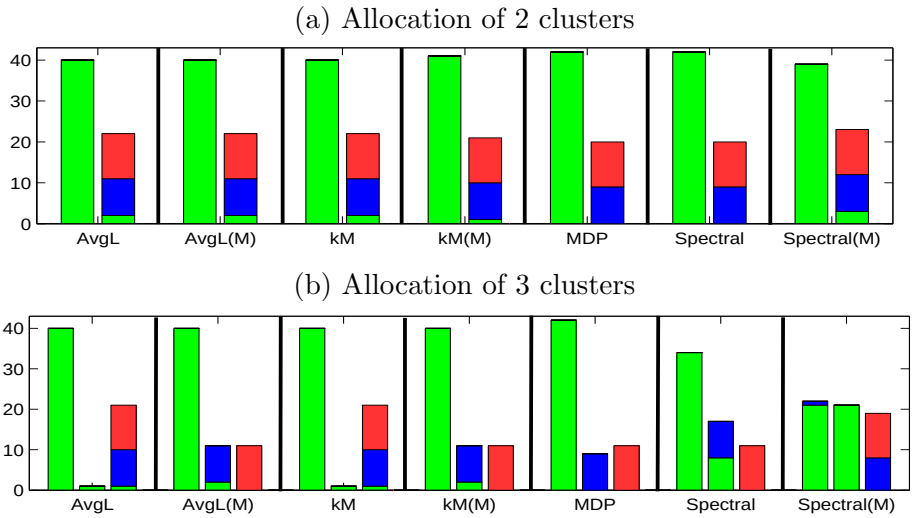

Figure 8: Compositions of (a) two and (b) three estimated clusters for Lymphoma data. Each bar corresponds to a single cluster, comprising of observations from different classes (DLBCL(■), FL(■) and CLL(■))

4.2 Synthetic control chart data

This data set contains 60 dimensional observations from 6 classes, viz., 'normal', 'cyclic', 'increasing trend', 'decreasing trend', 'upward shift' and 'downward shift'. There are 100 observations from each class. Figure 9 depicts a representation of the six classes.

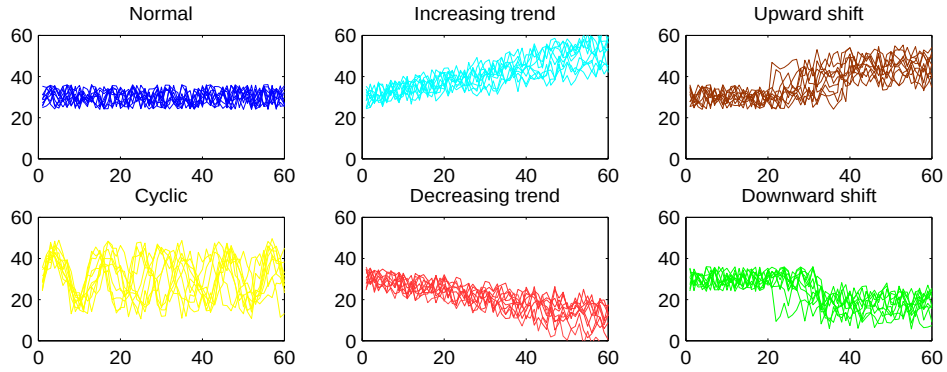

Figure 9: Six classes in 'Synthetic Control Chart' data

Methods based on Dunn index and PDI could find only two clusters in this data set when Euclidean distance was used. But they identified three clusters when their MADD versions were used, as did most other methods (see Table 4). Methods based on Gap

statistic again overestimated the number of clusters. JUMP statistic also overestimated the number of clusters in some cases. Though JUMP statistic based on MADD found six clusters in this data set when AvgL(M) was used for base clustering, those estimated clusters did not correspond to the six classes, as one can see from Figure 11.

Table 4: Number of clusters estimated by different methods in 'Control Chart' data

	Dunn	PDI	KL	GAP	JUMP	CV_a	CV_v
AvgL	2	2	3	10	8	3	3
AvgL(M)	3	3	10	8	6	3	3
kM	2	2	3	9	3	3	3
kM(M)	3	3	3	10	10	3	2

Since most of the methods identified two or three clusters in this data set, we used different clustering algorithms for estimating those two or three clusters. The results are shown as bar plots in Figure 10. Figure 10(a) shows that when different clustering algorithms were used to divide the data set into two groups, most of them put the observations from classes 'increasing trend' and 'upward shift' in one cluster and the rest in the other cluster. Only spectral clustering method based on Euclidean distance led to different cluster formations. MDP distance based clustering performed poorly in this data set; it formed one cluster by taking one or two observations from each class and put the rest of the observations in the other cluster. Note that, in this example, we had dimension smaller than sample size. So, poor performance of MDP clustering was expected in view of the results reported in Figure 2 and Table 1.

When we used the clustering algorithms to divide the data set into three clusters, average linkage and k-means algorithms (based on both Euclidean norm and MADD) formed one cluster consisting of observations from 'normal' and 'cyclic' classes; one cluster consisting of observations mainly from the classes 'decreasing trend' and 'downward shift', while the third cluster was formed mainly by observations from 'increasing trend' and 'upward shift' classes as before (see Figure 10(b)). Among these two types of algorithms, those based on average linkage had an edge. MDP distance based clustering performed poorly again and put almost all observations into a single cluster. Spectral clustering algorithm based on

(a) Allocation of 2 clusters

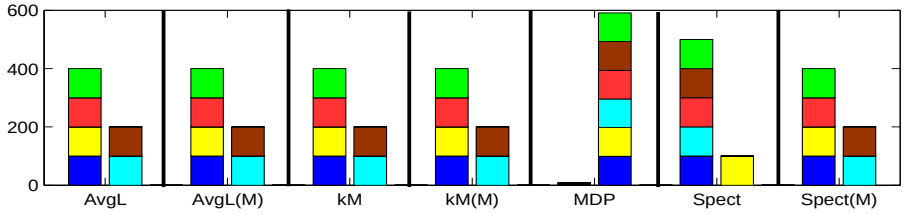

(b) Allocation of 3 clusters

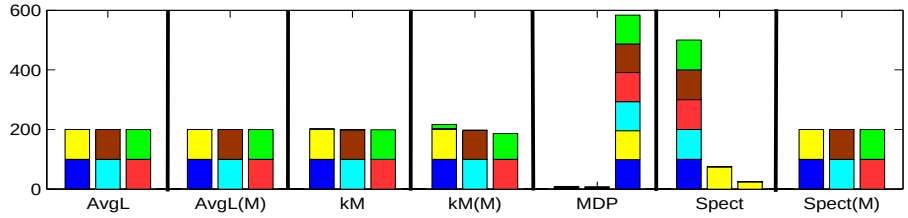

Figure 10: Compositions of (a) two and (b) three estimated clusters. Each bar corresponds to a single cluster, comprising of observations from different classes (normal(■), cyclic(■), increasing trend(■), decreasing trend(■), upward shift(■) and downward shift(■)).

Euclidean distance also performed poorly, but that based on MADD performed much better. It found the same clusters as obtained by average linkage methods.

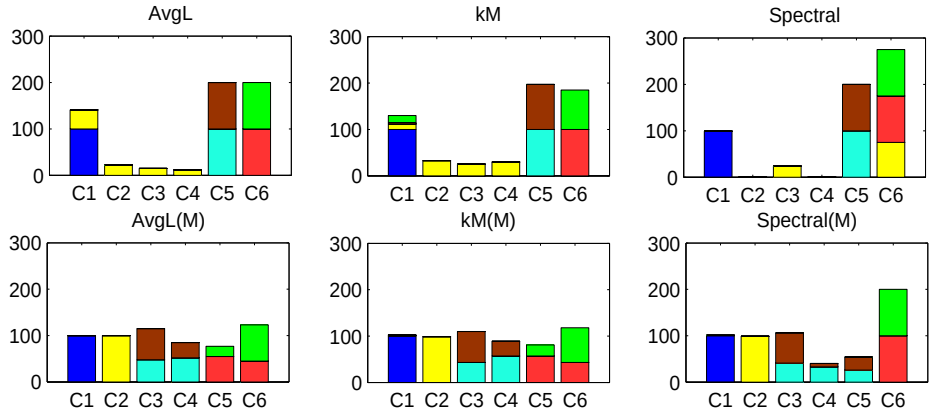

Figure 11: Compositions of six estimated clusters. Each bar corresponds to a single cluster, comprising of observations from different classes (normal(■), cyclic(■), increasing trend(■), decreasing trend(■), upward shift(■) and downward shift(■)).

We also tried to divide the data set into six clusters. In that situation, clustering algorithms based on MADD performed somewhat better than their Euclidean counterparts (see

Figure 11). For instance, while AvgL and kM put many observations from the 'cyclic' and 'normal' classes in the same cluster, AvgL(M) and kM(M) successfully separated normal, cyclic, upward and downward patterns. However, like all other clustering algorithms, they were unable to distinguish between 'increasing trend' and 'upward shift' or 'decreasing trend' and 'downward shift'. This was quite expected from the plot of the observations in Figure 9. MDP clustering put almost all observations into a single cluster as before. So we did not plot the results in Figure 11.

5 Concluding remarks

In high dimensions, concentration of Euclidean distance often has adverse effects on the performance of clustering algorithms based on it. In this article, we have used a data driven dissimilarity measure called MADD, which takes care of this problem. Clustering algorithms based on MADD can lead to perfect clustering for HDLSS data even when those based on Euclidean distance perform miserably. While MDP distance based clustering (see Ahn et al. 2012) performs poorly for dimension smaller than the sample size and even for HDLSS data with populations not differing in their locations, MADD based clustering algorithms can have excellent performance in those situations. We have amply demonstrated it using our theoretical results as well as analysis of several simulated and real data sets. In almost all the data sets, the spectral clustering algorithm of Shi and Malik (2000) also worked better when a MADD based similarity measure was used. We observed the same for the spectral clustering algorithm of Ng et al. (2002) as well, but to save space, we decided not to report them in this article.

Though we have defined MADD using the Euclidean distance, it can be defined using other distance measures as well. For instance, one can use l_p-distance with any $p \geq 1$ or other distance functions discussed in Biswas et al. (2015), and consistency of MADD can also be derived in those cases under appropriate regularity conditions. If the measurement vector contains some discrete/cateogorical variables, instead of Euclidean distance, we generally

consider other distance measures and k-medoids or other clustering algorithms are used instead of k-Means (see Hastie et al. 2009; Duda et al. 2012). MADD can also be defined based on those distance measures, and a MADD version of k-medoids algorithm can also be developed. Similarly, instead of AvgL(M), one can conveniently use MADD versions of other hierarchical methods.

We have also considered the problem of finding the number of clusters and we have seen that the methods based on Jump statistic and KL statistic usually perform better in high dimensions when their MADD versions are used. We have also successfully used MADD versions of Dunn index and penalized Dunn index for this purpose. Under appropriate regularity conditions, the methods based on KL statistic, Dunn index and penalized Dunn index turn out to be consistent in HDLSS asymptotic regime when AvgL(M) or kM(M) is used for base clustering. However, the choice of parameter λ in penalized Dunn index still remains an issue to be resolved. Though the consistency of our proposed method does not get affected by the choice of λ, its empirical performance may depend on it. Throughout this article, we have used $\lambda = 0.015$, which was chosen based on our empirical experience. But, a suitable data driven choice of λ may lead to further improvement.

Appendix: Proofs and mathematical details

Proof of Lemma 1: Since $\mathbf{X} \sim U_d(a_1, b_1)$, the distribution function of $R = d^{-1/2}\|\mathbf{X}\|$ is given by $F_R(r) = (r^d - a_1^d)/(b_1^d - a_1^d)$ for $a_1 \leq r \leq b_1$, 0 for $r < a_1$ and 1 for $r > B_1$. So, R has the density $f_R(r) = dr^{d-1}/(b_1^d - a_1^d)$ for $a_1 \leq r \leq b_1$ and 0 otherwise. Therefore,

$$E(d^{-1}\|\mathbf{X}\|^2) = E(R^2) = \frac{d}{d+2}\frac{b_1^{d+2} - a_1^{d+2}}{b_1^d - a_1^d} \to b_1^2, \text{ and}$$
$$E(d^{-2}\|\mathbf{X}\|^4) = E(R^4) = \frac{d}{d+4}\frac{b_1^{d+4} - a_1^{d+4}}{b_1^d - a_1^d} \to b_1^4 \tag{4}$$

as $d \to \infty$. This implies $Var(d^{-1}\|\mathbf{X}\|^2) \to 0$ and hence $d^{-1}\|\mathbf{X}\|^2 \xrightarrow{P} b_1^2$ as $d \to \infty$. Similarly, we have $d^{-1}\|\mathbf{Y}\|^2 \xrightarrow{P} b_2^2$ as $d \to \infty$.

Now, it is enough to show that $d^{-1}\langle \mathbf{X}, \mathbf{Y}\rangle \xrightarrow{P} 0$ as $d \to \infty$. Since $\mathbf{X}$ and $\mathbf{Y}$ are independent, and the components of $\mathbf{X}$ and $\mathbf{Y}$ are symmetric about zero, we have $E(d^{-1}\langle \mathbf{X}, \mathbf{Y}\rangle) = d^{-1}\sum_{i=1}^{d} E(X^{(i)})E(Y^{(i)}) = 0$. Also note that $\mathbf{X}$ and $\mathbf{Y}$ are spherically symmetric (see Fang et al. 1990). So, $E(X^{(i)}X^{(j)}) = E(Y^{(i)}Y^{(j)}) = 0$ for all $i \neq j$. Therefore, we have $Var(d^{-1}\langle \mathbf{X}, \mathbf{Y}\rangle) = E(d^{-2}\langle \mathbf{X}, \mathbf{Y}\rangle^2) = d^{-1}E(X^{(1)^2})E(Y^{(1)^2}) = d^{-1}E(d^{-1}\|\mathbf{X}\|^2)E(d^{-1}\|\mathbf{Y}\|^2)$. We have already proved that $E(d^{-1}\|\mathbf{X}\|^2) \to b_1^2$ and $E(d^{-1}\|\mathbf{Y}\|^2) \to b_2^2$ as $d \to \infty$. So, $Var(d^{-1}\langle \mathbf{X}, \mathbf{Y}\rangle) \to 0$ and hence $d^{-1}\langle \mathbf{X}, \mathbf{Y}\rangle \xrightarrow{P} 0$ as $d \to \infty$. □

Proof of Lemma 2: Symmetry and non-negativity of ρ is obvious. So, here we prove the triangle inequality. It is easy to see that

$$\begin{aligned}\big|\|\mathbf{x}_1 - \mathbf{x}_3\| - \|\mathbf{x}_2 - \mathbf{x}_3\|\big| &= \big|\|\mathbf{x}_1 - \mathbf{x}_2\| - \|\mathbf{x}_2 - \mathbf{x}_3\| - \|\mathbf{x}_1 - \mathbf{x}_2\| + \|\mathbf{x}_1 - \mathbf{x}_3\|\big| \\ &\leq \big|\|\mathbf{x}_1 - \mathbf{x}_2\| - \|\mathbf{x}_3 - \mathbf{x}_2\|\big| + \big|\|\mathbf{x}_2 - \mathbf{x}_1\| - \|\mathbf{x}_3 - \mathbf{x}_1\|\big|. \qquad (5)\end{aligned}$$

This proves the result for $n = 3$. If $n \geq 4$, for any $4 \leq k \leq n$, we have

$$\begin{aligned}\big|\|\mathbf{x}_1 - \mathbf{x}_k\| - \|\mathbf{x}_2 - \mathbf{x}_k\|\big| &= \big|\|\mathbf{x}_1 - \mathbf{x}_k\| - \|\mathbf{x}_3 - \mathbf{x}_k\| + \|\mathbf{x}_3 - \mathbf{x}_k\| - \|\mathbf{x}_2 - \mathbf{x}_k\|\big| \\ &\leq \big|\|\mathbf{x}_1 - \mathbf{x}_k\| - \|\mathbf{x}_3 - \mathbf{x}_k\|\big| + \big|\|\mathbf{x}_2 - \mathbf{x}_k\| - \|\mathbf{x}_3 - \mathbf{x}_k\|\big|. \qquad (6)\end{aligned}$$

Combining (5) and (6), we get

$$\sum_{k\neq 1,2}\big|\|\mathbf{x}_1 - \mathbf{x}_k\| - \|\mathbf{x}_2 - \mathbf{x}_k\|\big| \leq \sum_{k\neq 1,3}\big|\|\mathbf{x}_1 - \mathbf{x}_k\| - \|\mathbf{x}_3 - \mathbf{x}_k\|\big| + \sum_{k\neq 2,3}\big|\|\mathbf{x}_2 - \mathbf{x}_k\| - \|\mathbf{x}_3 - \mathbf{x}_k\|\big|$$

i.e., $\rho(\mathbf{x}_1, \mathbf{x}_2) \leq \rho(\mathbf{x}_1, \mathbf{x}_3) + \rho(\mathbf{x}_2, \mathbf{x}_3)$. □

Proof of Lemma 3: (i) Under $(A1)$ and $(A3)$, WLLN holds for the sequence $\{Z^{(i)};\ i \geq 1\}$, i.e., $\left|\frac{1}{d}\sum_{i=1}^{d} Z^{(i)} - \frac{1}{d}\sum_{i=1}^{d} E[Z^{(i)}]\right| \xrightarrow{P} 0$ as $d \to \infty$. Now, under $(A2)$, $\frac{1}{d}\sum_{i=1}^{d} E[Z^{(i)}]$ converges to $\sigma_k^2 + \sigma_{k'}^2 + \nu_{kk'}$ as $d \to \infty$.

(ii) Using (i), the result directly follows from the definition of ρ given in (2). □

Proof of Theorem 1: It follows from part (ii) of Lemma 3 that under $(A1) - (A3)$,

$$P\left(\max_{\mathbf{z},\mathbf{w}\in\mathcal{X}_\ell} \rho(\mathbf{z}, \mathbf{w}) < \min_{\mathbf{z}\in\mathcal{X}_i, \mathbf{w}\in\mathcal{X}_j} \rho(\mathbf{z}, \mathbf{w})\right) \to 1 \text{ as } d \to \infty, \qquad (7)$$

for every ℓ and $i \neq j$. So, at the first step of AvgL(M), two members of the same population will merge together with probability converging to 1 as $d \to \infty$. Now at any step r $(2 \leq r <$

$n - k_0$), given that observations from the same population were merged together at each of the $(r-1)$ previous steps, any cluster C becomes a subset of $\mathcal{X}_j$ for some j, and we have

$$P\left(\max_{\ell} \max_{C,C' \subset \mathcal{X}_\ell} \Delta(C,C') < \min_{i \neq j} \min_{C \subset \mathcal{X}_i, C' \subset \mathcal{X}_j} \Delta(C,C') \right) \to 1 \text{ as } d \to \infty, \tag{8}$$

where $\Delta(C,C') = \frac{1}{|C||C'|} \sum_{\mathbf{Z} \in C, \mathbf{W} \in C'} \rho(\mathbf{Z},\mathbf{W})$. Therefore, two clusters containing observations from the same population will merge with probability tending to 1 as $d \to \infty$. Since k_0 is known, these two facts together prove the result. □

Proof of Theorem 2: Note that for any k, $d^{-1}|C_k|^{-1} \sum_{\mathbf{Z} \in C_k, \mathbf{Z} \neq \mathbf{X}} \rho^2(\mathbf{X},\mathbf{Z}) \xrightarrow{P} 0$ as $d \to \infty$ if and only if $\mathbf{x}$ and all the observations in C_k are from the same population (follows from Lemma 3(ii)). So, if each C_k $(k = 1, \ldots, k_0)$ contains observations from the same population, $\varphi^*(C_1, \ldots, C_{k_0}) \xrightarrow{P} 0$ as $d \to \infty$. Otherwise, $\varphi^*(C_1, \ldots, C_{k_0})$ converges (in probability) to a positive constant. So, when k_0 is known, for the minimization of $\varphi^*(C_1, \ldots, C_{k_0})$, each C_k must contain all observations from a single population with probability converging to one as the dimension increases. This proves the convergence of the Rand index to zero. □

We now state a lemma, which will be helpful in proving Theorems 3 and 4.

Lemma 4. *Suppose that we have n independent observations from k_0 populations satisfying $(A1), (A2)$ and $(A3^\circ)$. Then, for any two observations $\mathbf{X}$ and $\mathbf{X}'$ from the same population, $\rho(\mathbf{X},\mathbf{X}')$ is of the order $\boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$, for some $0 < \varepsilon_0 < 1/2$.*

Proof of Lemma 4: Suppose that $\mathbf{X}$ and $\mathbf{X}'$ are from the k-th population and $\mathbf{Y}$ is from the k'-th population $(1 \leq k, k' \leq k_0)$. To prove the result, it is enough to show that $\|\mathbf{X}-\mathbf{Y}\| - \|\mathbf{X}'-\mathbf{Y}\| = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$. First assume that $\varepsilon < 1/2$. Let us define $V_d = \|\mathbf{X}-\mathbf{Y}\|^2$ and $V_d' = \|\mathbf{X}'-\mathbf{Y}\|^2$. Using first order Taylor's series expansion of $f(x) = \sqrt{x}$, we get

$$\sqrt{V_d} = \sqrt{E(V_d)} + \frac{V_d - E(V_d)}{2\sqrt{V_d^*}} \text{ and } \sqrt{V_d'} = \sqrt{E(V_d')} + \frac{V_d' - E(V_d')}{2\sqrt{V_d'^*}},$$

where $|V_d^* - E(V_d)| < |V_d - E(V_d)|$ and $|V_d'^* - E(V_d')| < |V_d' - E(V_d')|$. This implies

$$\begin{aligned} \frac{\|\mathbf{X}-\mathbf{Y}\| - \|\mathbf{X}'-\mathbf{Y}\|}{d^{\frac{1}{2}-\varepsilon}} &= \frac{\sqrt{V_d} - \sqrt{V_d'}}{d^{\frac{1}{2}-\varepsilon}} \\ &= \frac{\sqrt{E(V_d)} - \sqrt{E(V_d')}}{d^{\frac{1}{2}-\varepsilon}} + \frac{V_d - E(V_d)}{2d^{\frac{1}{2}-\varepsilon}\sqrt{V_d^*}} - \frac{V_d' - E(V_d')}{2d^{\frac{1}{2}-\varepsilon}\sqrt{V_d'^*}}. \end{aligned} \tag{9}$$

Since $E(V_d) = E(V_d')$, the first term on the right hand side is zero. Note that the second term can be written as

$$\frac{V_d - E(V_d)}{2d^{\frac{1}{2}-\varepsilon}\sqrt{V_d^*}} = \frac{V_d - E(V_d)}{\sqrt{4Var(V_d)}}\sqrt{\frac{Var(V_d)}{d^{2(1-\varepsilon)}}}\sqrt{\frac{d}{V_d^*}}$$

Using Markov's inequality, it is easy to check that $[V_d - E(V_d)]/\sqrt{4Var(V_d)} = \boldsymbol{O}_P(1)$. Under $(A1), (A2)$ and $(A3^\circ)$, we also have (a) $Var(V_d)/d^{2(1-\varepsilon)} = \boldsymbol{O}(1)$ and (b) $d^{-1}|V_d^* - E(V_d)| \leq d^{-1}|V_d - E(V_d)| \xrightarrow{P} 0$ as $d \to \infty$. Since $d^{-1}E(V_d) \to \sigma_k^2 + \sigma_{k'}^2 + \nu_{kk'}$ as $d \to \infty$, from (b), we have $\sqrt{d/V_d^*} \xrightarrow{p} (\sigma_k^2 + \sigma_{k'}^2 + \nu_{kk'})^{-1/2}$. Thus, we get $\frac{V_d - E(V_d)}{d^{\frac{1}{2}-\varepsilon}\sqrt{V_d^*}} = \boldsymbol{O}_P(1)$.

Similarly, one can show that $\frac{V_d' - E(V_d')}{d^{\frac{1}{2}-\varepsilon}\sqrt{V_d'^*}} = \boldsymbol{O}_P(1)$. Combining all these results, from (9) we get $\|\mathbf{X} - \mathbf{Y}\| - \|\mathbf{X}' - \mathbf{Y}\| = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon})$, which shows that the lemma holds for $0 < \varepsilon < 1/2$ with $\varepsilon_0 = \varepsilon$. For $\varepsilon > 1/2$, one can easily check that the above arguments hold for any $0 < \varepsilon_0 < 1/2$, for instance, $\varepsilon_0 = 1/4$. This proves the lemma. □

All the estimation methods that we discuss henceforth are based on MADD. So, $\Delta_0(C_i) = \frac{1}{|C_i|(|C_i|-1)}\sum_{\mathbf{z},\mathbf{w}\in C_i}\rho(\mathbf{z},\mathbf{w})$ for every i, $\Delta(C_i, C_j) = \frac{1}{|C_i||C_j|}\sum_{\mathbf{z}\in C_i,\mathbf{w}\in C_j}\rho(\mathbf{z},\mathbf{w})$ for every $i \neq j$ and $W_k = \sum_{j=1}^{k}\frac{1}{2|C_j|}\sum_{\mathbf{z},\mathbf{w}\in C_j}\rho^2(\mathbf{z},\mathbf{w})$ for every $k \geq 1$.

Proof of Theorem 3: (i) Since the base clustering algorithm is POP at k_0, for any $k < k_0$, there exists at least one estimated cluster C_i which contains observations from two different populations, and no two different clusters contain observations from the same population. So, we have $\Delta_0(C_i) = \boldsymbol{O}_P(d^{\frac{1}{2}})$ for some i and $\Delta(C_i, C_j) = \boldsymbol{O}_P(d^{\frac{1}{2}})$ for every $i \neq j$ (follows from Lemma 3(ii)). Thus, $B_k^\circ = \min_{1\leq i<j\leq k}\Delta(C_i, C_j) = \boldsymbol{O}_P(d^{\frac{1}{2}})$ and $W_k^\circ = \max_{1\leq i\leq k}\Delta_0(C_i) = \boldsymbol{O}_P(d^{\frac{1}{2}})$, and this implies $\mathrm{DI}(k) = \boldsymbol{O}_P(1)$.

For $k > k_0$, no cluster contains observations from two different populations, while there exists at least two clusters which contain observations from the same population. So, for some $0 < \varepsilon_0 < 1/2$, $\Delta_0(C_i) = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ for every i and $\Delta(C_i, C_j) = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ for some $i \neq j$ (follows from Lemma 4). Thus, we have $B_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ and $W_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$, and hence $\mathrm{DI}(k) = \boldsymbol{O}_P(1)$.

For $k = k_0$, each cluster contains observations from same population and two different clusters contain observations from two different populations. This implies that $\Delta_0(C_i) = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ for every i and $\Delta(C_i, C_j) = \boldsymbol{O}_P(d^{\frac{1}{2}})$ for every $i \neq j$. So, we have $B_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}})$ and $W_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ and hence $\mathrm{DI}(k) = \boldsymbol{O}_P(d^{\varepsilon_0})$.

Combining these three cases, we get $P\Big(\mathrm{DI}(k_0) > \mathrm{DI}(k) \quad \forall k \neq k_0\Big) \to 1$ as $d \to \infty$. This implies $\hat{k}_D^* \xrightarrow{P} k_0$ as $d \to \infty$.

(ii) Note that for all $k \geq 2$, the KL statistic can be written as $\mathrm{KL}(k) = \left|\frac{\Lambda_{k-1}-\Lambda_k}{\Lambda_k-\Lambda_{k+1}}\right|$, where $\Lambda_k = k^{2/d} W_k$. Since the base clustering algorithm is POP at k_0, from our discussion in part (i), it follows that $W_k = \boldsymbol{O}_P(d)$ for $k < k_0$ and $W_k = \boldsymbol{O}_P(d^{1-2\varepsilon_0})$ for $k \geq k_0$. Also, note that for any fixed k, $k^{2/d} \to 1$ as $d \to \infty$. So, for all $k \geq 2$, Λ_k has the same asymptotic order as W_k. Now, it is easy to check that $\mathrm{KL}(k_0)$ is of the order $\boldsymbol{O}_P(d^{1-2\varepsilon_0})$ and $\mathrm{KL}(k) = \boldsymbol{O}_P(1)$ for all other values of k. Therefore, the probability $P\Big(\mathrm{KL}(k_0) > \mathrm{KL}(k) \quad \forall k \neq k_0\Big) \to 1$ as $d \to \infty$, and hence we have $P(\hat{k}_{KL}^* = k_0) \to 1$ as $d \to \infty$.

(iii) Note that $\hat{d}_k = W_k/d$ is of the order $\boldsymbol{O}_P(1)$ and $\boldsymbol{O}_P(d^{-2\varepsilon_0})$ for $1 \leq k < k_0$ and $k \geq k_0$, respectively (follows from our discussion in parts (i) and (ii)). So, we have $\textsc{Jump}(k) = \hat{d}_k^{-1} - \hat{d}_{k-1}^{-1} = \boldsymbol{O}_P(1)$ or $\boldsymbol{O}_P(d^{2\varepsilon_0})$ according as $k < k_0$ or $k \geq k_0$. This implies that $P\big(\hat{k}_J^* < k_0\big) \to 0$ as $d \to \infty$. □

Proof of Theorem 4: First consider the case $k_0 \geq 2$. While proving part (i) of Theorem 3, we have shown that in this case, $\mathrm{DI}(k) = \boldsymbol{O}_P(1)$ for $2 \leq k \neq k_0$ and $\mathrm{DI}(k_0) = \boldsymbol{O}_P(d^{\varepsilon_0})$ for some $\varepsilon_0 > 0$. From Lemma 3(ii), it is also evident that $W_1^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}})$ and $B_1^\circ \overset{def}{=} B_2^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}})$, thus $\mathrm{DI}(1) \overset{def}{=} B_1^\circ/W_1^\circ = \boldsymbol{O}_P(1)$. So, for any $k \neq k_0$, we have $\mathrm{PDI}_\lambda(k_0) - \mathrm{PDI}_\lambda(k) = \mathrm{DI}(k_0) - \mathrm{DI}(k) - \lambda(k_0 - k)\log(d) \to \infty$ as $d \to \infty$. This implies $\hat{k}_{PD}^* \xrightarrow{P} k_0$ as $d \to \infty$.

Now when $k_0 = 1$, for every $k \geq 1$ we have, $W_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$, $B_k^\circ = \boldsymbol{O}_P(d^{\frac{1}{2}-\varepsilon_0})$ (follows from Lemma 4). Thus, $\mathrm{PDI}_\lambda(1) - \mathrm{PDI}_\lambda(k) = \mathrm{DI}(1) - \mathrm{DI}(k) + (k-1)\lambda\log(d) \xrightarrow{P} \infty$ as $d \to \infty$. This proves the consistency of $\hat{k}_{PD}^*$. □

References

Ahn, J., Lee, M. H. and Yoon, Y. J. (2012) Clustering high dimension, low sample size data using the maximal data piling distance. *Statist. Sinica*, **22**, 443–464.

Alizadeh, A. A., Eisen, M. B., Davis, R. E., Ma, C., Lossos, I. S., Rosenwald, A., Boldrick, J. C., Sabet, H., Tran, T., Yu, X., Powell, J. I., Yang, L., Marti, G. E., Moore, T., Hudson, James, J., Lu, L., Lewis, D. B., Tibshirani, R., Sherlock, G., Chan, W. C., Greiner, T. C., Weisenburger, D. D., Armitage, J. O., Warnke, R., Levy, R., Wilson, W., Grever, M. R., Byrd, J. C., Botstein, D., Brown, P. O. and Staudt, L. M. (2000) Distinct types of diffuse large B-cell lymphoma identified by gene expression profiling. *Nature*, **403**, 503–511.

Biswas, M., Mukhopadhyay, M. and Ghosh, A. K. (2014) A distribution-free two-sample run test applicable to high-dimensional data. *Biometrika*, **101**, 913–926.

— (2015) On some exact distribution-free one-sample tests for high dimension low sample size data. *Statist. Sinica*, **25**, 1421–1435.

Calinski, R. B. and Harabasz, J. (1974) A dendrite method of cluster analysis. *Comm. Statist.*, **3**, 1–27.

Duda, R. O., Hart, P. E. and Stork, D. G. (2012) *Pattern classification.* Wiley, New York.

Dunn, J. C. (1973) A fuzzy relative of the isodata process and its use in detecting compact well-separated clusters. *J. Cybern.*, **3**, 32–57.

Fang, K. T., Kotz, S. and Ng, K. W. (1990) *Symmetric multivariate and related distributions.* Chapman & Hall, London.

Hall, P., Marron, J. S. and Neeman, A. (2005) Geometric representation of high dimension, low sample size data. *J. Royal Statist. Soc. Ser. B*, **67**, 427–444.

Hartigan, J. A. (1975) *Cluster Algorithms.* Wiley, New York.

Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference and Prediction.* Springer, New York, 2nd edn.

Kaufman, L. and Rousseeuw, P. (1990) *Finding Groups in Data: An Introduction to Cluster Analysis.* Wiley, New York.

Krzanowski, W. J. and Lai, Y. T. (1985) A criterion for determining the number of clusters in a data set. *Biometrics*, **44**, 23–34.

von Luxburg, U. (2007) A tutorial on spectral clustering. *Statistics and Computing*, **17**, 395–416.

Ng, A. Y., Jordan, M. I. and Weiss, Y. (2002) On spectral clustering: Analysis and an algorithm. *Adv. in Neural Inf. Proc. Syst.*, **2**, 849–856.

Rand, W. M. (1971) Objective criteria for the evaluation of clustering methods. *J. Amer. Statist. Assoc.*, **66**, 846–850.

Shi, J. and Malik, J. (2000) Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, **22**, 888–905.

Sugar, C. A. and James, G. M. (2003) Finding the number of clusters in a dataset. *J. Amer. Statist. Assoc.*, **98**, 750–763.

Tibshirani, R., Walther, G. and Hastie, T. (2001) Estimating the number of clusters in a data set via the gap statistic. *J. Royal Statist. Soc. Ser. B*, **63**, 411–423.

Wang, J. (2010) Consistent selection of the number of clusters via crossvalidation. *Biometrika*, **97**, 893–904.